arXiv:1503.02164v1 [cs.IT] 7 Mar 2015

A Nonconvex Approach for Structured Sparse Learning

Shubao Zhang Hui Qian
College of Computer Science and Technology
Zhejiang University, Hangzhou 310027, China
{bravemind, qianhui}@zju.edu.cn
Zhihua Zhang
College of Computer Science and Engineering
Shanghai Jiao Tong University, Shanghai 200240, China
zhihua@sjtu.edu.cn

Abstract

Sparse learning is an important topic in many areas such as machine learning, statistical estimation, signal processing, etc. Recently, there emerges a growing interest on structured sparse learning. In this paper we focus on the ℓ_q-analysis optimization problem for structured sparse learning ($0 < q \leq 1$). Compared to previous work, we establish weaker conditions for exact recovery in noiseless case and a tighter non-asymptotic upper bound of estimate error in noisy case. We further prove that the nonconvex ℓ_q-analysis optimization can do recovery with a lower sample complexity and in a wider range of cosparsity than its convex counterpart. In addition, we develop an iteratively reweighted method to solve the optimization problem under the variational framework. Theoretical analysis shows that our method is capable of pursuing a local minima close to the global minima. Also, empirical results of preliminary computational experiments illustrate that our nonconvex method outperforms both its convex counterpart and other state-of-the-art methods.

1 Introduction

The sparse learning problem is widely studied in many areas including machine learning, statistical estimate, compressed sensing, image processing and signal processing, etc. Typically, this problem can be defined as the following linear model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{w}, \tag{1}$$

where $\boldsymbol{\beta} \in \mathbb{R}^d$ is the vector of regression coefficients, $\mathbf{X} \in \mathbb{R}^{m\times d}$ is a design matrix with possibly far fewer rows than columns, $\mathbf{w} \in \mathbb{R}^m$ is a noise vector, and $\mathbf{y} \in \mathbb{R}^m$ is the noisy observation. As is well known, learning with the

ℓ_1 norm (convex relaxation of the ℓ_0 norm), such as lasso [Tibshirani, 1996] or basis pursuit [Chen et al., 1998], encourages sparse estimate of $\boldsymbol{\beta}$. Recently, this approach has been extended to define structured sparsity. Tibshirani and Taylor [2011] proposed the generalized lasso

$$\min_{\boldsymbol{\beta}} \frac{1}{2}||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}||_2^2 + \lambda||\mathbf{D}\boldsymbol{\beta}||_1, \tag{2}$$

which assumes that the parameter $\boldsymbol{\beta}$ is sparse under a linear transformation $\mathbf{D} \in \mathbb{R}^{n\times d}$. An equivalent constrained version is the ℓ_1-analysis minimization proposed by Candés et al. [2010], i.e.,

$$\min_{\boldsymbol{\beta}} ||\mathbf{D}\boldsymbol{\beta}||_1 \quad \text{s.t.} \quad ||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}||_2 \leq \epsilon, \tag{3}$$

where $\mathbf{D}$ is called the analysis operator. In contrast to the lasso and basis pursuit in $\mathbf{D} = \mathbf{I}$, the generalized lasso and ℓ_1-analysis minimization make a structured sparsity assumption so that it can explore structures on the parameter. They include several well-known models as special cases, e.g., fused lasso [Tibshirani et al., 2005], generalized fused lasso [Viallon et al., 2014], edge Lasso [Sharpnack et al., 2012], total variation (TV) minimization [Rudin et al., 1992], trend filtering [Kim et al., 2009], the LLT model [Lysaker et al., 2003], the inf-convolution model [Chambolle and Lions, 1997], etc. Additionally, the generalized lasso and ℓ_1-analysis minimization have been demonstrated to be effective and even superior over the standard sparse learning in many application problems.

The seminal work of Fan and Li [2001] showed that the nonconvex sparse learning holds better properties than the convex one. Motivated by that, this paper investigates the following ℓ_q-analysis minimization ($0 < q \leq 1$) problem

$$\min_{\boldsymbol{\beta}} ||\mathbf{D}\boldsymbol{\beta}||_q^q \quad \text{s.t.} \quad ||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}||_2 \leq \epsilon. \tag{4}$$

We consider both theoretical and computational aspects. We summary the major contributions as follows:

- We establish weaker conditions for exact recovery in noiseless case and a tighter non-asymptotic upper bound of estimate error in noisy case. Particularly, we provide a necessary and sufficient condition guaranteeing exact recovery via the ℓ_q-analysis minimization. To the best of our knowledge, our work is the first study in this issue.

- We show the advantage of the nonconvex ℓ_q-analysis minimization ($q < 1$) over its convex counterpart. Specifically, the nonconvex ℓ_q-analysis minimization can do recovery with a lower sample complexity (on the order of $qk\log(n/k)$) and in a wider range of cosparsity.

- We resort to an iteratively reweighted method to solve the ℓ_q-analysis minimization problem. Furthermore, we prove that our method is capable to obtain a local minima close to the global minima.

The numerical results are consistent with the theoretical analysis. For example, the nonconvex ℓ_q-analysis minimization indeed can do recovery with smaller sample size and in a wider range of cosparsity than the convex method. The numerical results also show that our iteratively reweighted method outperforms the other state-of-the-art methods such as NESTA [Becker et al., 2011], split Bregman method [Cai et al., 2009a], and iteratively reweighted ℓ_1 method [Candés et al., 2007] for the ℓ_1-analysis minimization problem and the greedy analysis pursuit (GAP) method [Nam et al., 2011] for the ℓ_0-analysis minimization problem ($q \to 0$ in (4)) .

1.1 Related Work

Candés et al. [2010] studied the ℓ_1-analysis minimization problem in the setting that the observation is contaminated with stochastic noise and the analysis vector $\mathbf{D}\boldsymbol{\beta}$ is approximately sparse. They provided a ℓ_2 norm estimate error bounded by $C_0\epsilon + C_1 k^{-1/2}||\mathbf{D}\boldsymbol{\beta} - (\mathbf{D}\boldsymbol{\beta})(k)||_1$ under the assumption that $\mathbf{X}$ obeys the D-RIP condition $\delta_{2k} < 0.08$ or $\delta_{7k} < 0.6$ and $\mathbf{D}$ is a Parseval tight frame [1]. Nam et al. [2011] studied the ℓ_1-analysis minimization problem in the setting that there is no noise and the analysis vector $\mathbf{D}\boldsymbol{\beta}$ is sparse. They showed that a null space property with sign pattern is necessary and sufficient to guarantee exact recovery. Liu et al. [2012] improved the analysis in [Candés et al., 2010]. They established an estimate error bound similar to the one in [Candés et al., 2010] for the general frame case. And for the Parseval frame case, they provided a weaker D-RIP condition $\delta_{2k} < 0.2$.

Tibshirani and Taylor [2011] proposed the generalized lasso and developed a LARS-like algorithm pursuing its solution path. Vaiter et al. [2013] conducted a robustness analysis of the generalized lasso against noise. Liu et al. [2013] derived an estimate error bound for the generalized lasso under the assumption that the condition number of $\mathbf{D}$ is bounded. Specifically, a ℓ_2 norm estimate error bounded by $C\lambda + ||(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{w}||_2$ is provided. Needell and Ward [2013] investigated the total variation minimization. They proved that for an image $\boldsymbol{\beta} \in \mathbb{R}^{N\times N}$, the TV minimization can stably recover it with estimate error less than $C\log(\frac{N^2}{k})(\epsilon + ||\mathbf{D}\boldsymbol{\beta} - (\mathbf{D}\boldsymbol{\beta})(k)||_1/\sqrt{k})$ when the sampling matrix satisfies the RIP of order k.

So far, all the related works discussed above consider convex optimization problem. Aldroubi et al. [2012] first studied the nonconvex ℓ_q-analysis minimization problem (4). They established estimate error bound using the null space property and restricted isometry property respectively. For the Parseval frame case, they showed that the D-RIP condition $\delta_{7k} < \frac{6-3(2/3)^{2/q-2}}{6-(2/3)^{2/q-2}}$ is sufficient to guarantee stable recovery. Li and Lin [2014] showed that the D-RIP condition $\delta_{2k} < 0.5$ is sufficient to guarantee the success of ℓ_q-analysis minimization. In this paper, we significantly improve the analysis of ℓ_q-analysis

[1] A set of vectors $\{\mathbf{d}_k\}$ is a frame of $\mathbb{R}^d$ if there exist constants $0 < A \le B < \infty$ such that $\forall \boldsymbol{\beta} \in \mathbb{R}^d$, $A||\boldsymbol{\beta}||_2^2 \le ||D\boldsymbol{\beta}||_2^2 \le B||\boldsymbol{\beta}||_2^2$, where $\{\mathbf{d}_k\}$ are the columns of $\mathbf{D}^T$. When $A = B = 1$, the columns of $\mathbf{D}^T$ form a Parseval tight frame and $\mathbf{D}^T\mathbf{D} = \boldsymbol{I}$.

minimization. For example, we provide a weaker D-RIP condition $\delta_{2k} < \frac{\sqrt{2}}{2}$. Additionally, we show the advantage of the nonconvex ℓ_q-analysis minimization over its convex counterpart.

2 Preliminaries

Throughout this paper, $\mathbb{N}$ denotes the natural number. $\lfloor \cdot \rfloor$ denotes the rounding down operator. The i-th entry of a vector $\boldsymbol{\beta}$ is denoted by β_i. The *best k-term approximation* of a vector $\boldsymbol{\beta} \in \mathbb{R}^d$ is obtained by setting its $d-k$ insignificant components to zero and denoted by $\boldsymbol{\beta}(k)$. The ℓ_q norm of a vector $\boldsymbol{\beta} \in \mathbb{R}^d$ is defined as $||\boldsymbol{\beta}||_q = (\sum_{i=1}^d |\beta_i|^q)^{1/q}$ [2] for $0 < q < \infty$. When q tends to zero, $||\boldsymbol{\beta}||_q^q$ is the ℓ_0 norm $||\boldsymbol{\beta}||_0$ used to measure the *sparsity* of $\boldsymbol{\beta}$. $\sigma_k(\boldsymbol{\beta})_q = \inf_{\mathbf{z}\in\{\mathbf{z}\in\mathbb{R}^d:||\mathbf{z}||_0\leq k\}} ||\boldsymbol{\beta}-\mathbf{z}||_q$ denotes the best k-term approximation error of $\boldsymbol{\beta}$ with the ℓ_q norm. The i-th row of a matrix $\mathbf{D}$ is denoted by $\mathbf{D}_{i.}$. $\sigma_{max}(\mathbf{D})$ and $\sigma_{min}(\mathbf{D})$ denote the maximal and minimal nonzero singular value of $\mathbf{D}$, respectively. Let $\kappa = \frac{\sigma_{max}(\mathbf{D})}{\sigma_{min}(\mathbf{D})}$, and Null($\mathbf{X}$) denote the null space of $\mathbf{X}$.

Now we introduce some concepts related to the ℓ_q-analysis minimization problem (4). The number of zeros in the analysis vector $\mathbf{D}\boldsymbol{\beta}$ is refered to as *cosparsity* [Nam et al., 2011], and defined as $l := n - ||\mathbf{D}\boldsymbol{\beta}||_0$. Such a vector $\boldsymbol{\beta}$ is said to be l-cosparse. The *support* of a vector $\boldsymbol{\beta}$ is the collection of indices of nonzeros in the vector, denoted by $T := \{i : \beta_i \neq 0\}$. T^c denotes the complement of T. The indices of zeros in the analysis vector $\mathbf{D}\boldsymbol{\beta}$ is defined as the *cosupport* of $\boldsymbol{\beta}$, and denoted by $\Lambda := \{j : \langle \mathbf{D}_{j.}, \boldsymbol{\beta}\rangle = 0\}$. The submatrix $\mathbf{D}_T$ is constructed by replacing the rows of $\mathbf{D}$ corresponding to T^c by zero rows. Denote $\mathbf{D}_T\boldsymbol{\beta} = (\mathbf{D}\boldsymbol{\beta})_T$. Based on these concepts, we can see that a l-cosparse vector $\boldsymbol{\beta}$ lies in the subspace $\mathcal{W}_\Lambda := \{\boldsymbol{\beta} : \mathbf{D}_\Lambda\boldsymbol{\beta} = \mathbf{0}, |\Lambda| = l\} = \text{Null}(\mathbf{D}_\Lambda)$. Here $|\Lambda|$ is the cardinality of Λ.

In our analysis below, we use the notion of $\mathcal{A}$-RIP [Blumensath and Davies, 2008].

Definition 1 ($\mathcal{A}$-restricted isometry property) *A matrix $\boldsymbol{\Phi} \in \mathbb{R}^{m\times d}$ obeys the $\mathcal{A}$-restricted isometry property with constant $\delta_\mathcal{A}$ over any subset $\mathcal{A} \in \mathbb{R}^d$, if $\delta_\mathcal{A}$ is the smallest quantity satisfying*

$$(1-\delta_\mathcal{A})||\mathbf{v}||_2^2 \leq ||\boldsymbol{\Phi}\mathbf{v}||_2^2 \leq (1+\delta_\mathcal{A})||\mathbf{v}||_2^2$$

for all $\mathbf{v} \in \mathcal{A}$.

Note that RIP [Candés and Tao, 2004], D-RIP [Candés et al., 2010] and $\boldsymbol{\Omega}$-RIP [Giryes et al., 2013] are special instances of the $\mathcal{A}$-RIP with different choices of the set $\mathcal{A}$. For example, when choosing $\mathcal{A} = \{\mathbf{D}\mathbf{v} : \mathbf{v} \in \mathbb{R}^d, ||\mathbf{v}||_0 \leq k\}$ and $\mathcal{A} = \{\mathbf{v} : \mathbf{v} \in \mathbb{R}^d, \mathbf{D}_\Lambda\mathbf{v} = \mathbf{0}, |\Lambda| \geq l\}$, the corresponding $\mathcal{A}$-restricted isometries are D-RIP and Ω-RIP, respectively. It has been verified that any random matrix $\boldsymbol{\Phi}$ holds the $\mathcal{A}$-restricted isometry property with overwhelming probability provided that the number of samples depends logarithmically on the number of subspaces in $\mathcal{A}$ [Blumensath and Davies, 2008].

[2] $||\boldsymbol{\beta}||_q$ for $0 < q < 1$ is not a norm, but $d(\mathbf{u},\mathbf{v}) = ||\mathbf{u}-\mathbf{v}||_q^q$ for $\mathbf{u},\mathbf{v} \in \mathbb{R}^d$ is a metric.

3 Main Results

In this section, we present our main theoretical results pertaining to the ability of ℓ_q-analysis minimization to estimate (approximately) cosparse vectors with and without noise.

3.1 Exact Recovery in Noiseless Case

A well-known necessary and sufficient condition guaranteeing the success of basis pursuit is the null space property [Cohen et al., 2009]. Naturally, we define a null space property adapted to $\mathbf{D}$ (D-NSPq) of order k [Aldroubi et al., 2012] for the ℓ_q-analysis minimization. That is,

$$\forall \boldsymbol{v} \in \mathrm{Null}(\mathbf{X})/\{\mathbf{0}\}, \forall |T| \leq k, ||\boldsymbol{D}_T \boldsymbol{v}||_q^q < ||\boldsymbol{D}_{T^c} \boldsymbol{v}||_q^q. \tag{5}$$

Theorem 1 *Let $\boldsymbol{\beta} \in \mathbb{R}^d$ with cosupport Λ, $||\mathbf{D}\boldsymbol{\beta}||_0 = k$, and $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$. Then $\boldsymbol{\beta}$ is the unique minimizer of the ℓ_q-analysis minimization (4) with $\epsilon = 0$ if and only if $\mathbf{X}$ satisfies the D-NSPq (5) relative to Λ^c.*

Letting the set Λ ($|\Lambda| = l$) vary, the following result is a corollary of Theorem 1.

Corollary 1 *Given a matrix $\mathbf{X} \in \mathbb{R}^{m \times d}$ and $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$, the ℓ_q-analysis minimization (4) with $\epsilon = 0$ recovers every l-cosparse vector $\boldsymbol{\beta} \in \mathbb{R}^d$ as a unique minimizer if and only if $\mathbf{X}$ satisfies the D-NSPq (5) of order $n - l$.*

This corollary establishes a necessary and sufficient condition for exact recovery of all l-cosparse vectors via the ℓ_q-analysis minimization. It also implies that for every $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$ with l-cosparse $\boldsymbol{\beta}$, the ℓ_q-analysis minimization actually solves the ℓ_0-analysis minimization when the D-NSPq of order $n-l$ holds. Based on the D-NSPq, the following corollary shows that the nonconvex ℓ_q-analysis minimization is not worse than its convex counterpart.

Corollary 2 *For $0 < q_1 < q_2 \leq 1$, the sufficient condition for exact recovery via the ℓ_{q_2}-analysis minimization is also sufficient for exact recovery via the ℓ_{q_1}-analysis minimization.*

It is hard to check the D-NSPq (5). The following theorem provides a sufficient condition for exact recovery using the $\mathcal{A}$-RIP.

Theorem 2 *Let $\boldsymbol{\beta} \in \mathbb{R}^d$, $||\mathbf{D}\boldsymbol{\beta}||_0 = k$, and $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$. Assume that $\mathbf{D} \in \mathbb{R}^{n \times d}$ has full column rank, and its condition number is upper bounded by $\kappa < \sqrt{\frac{2\rho+1+\sqrt{4\rho+1}}{2\rho}}$. If $\mathbf{X} \in \mathbb{R}^{m \times d}$ satisfies the $\mathcal{A}$-RIP over the set $\mathcal{A} = \{\mathbf{Dv} : ||\mathbf{v}||_0 \leq (t^q+1)k\}$ with $k, t^q k \in \mathbb{N}, t > 0, q \in (0, 1]$, i.e.,*

$$\delta_{(t^q+1)k} < \frac{\rho(1-\kappa^4) + \kappa^2\sqrt{4\rho+1}}{\rho(\kappa^2+1)^2 + \kappa^2} \tag{6}$$

with $\rho = t^{q-2}/4$, then $\boldsymbol{\beta}$ is the unique minimizer of the ℓ_q-analysis minimization (4) with $\epsilon = 0$.

Table 1: Different Sufficient Conditions

q	t	κ	Recovery condition
1	1	1	$\delta_{2k} < \frac{\sqrt{2}}{2}$
$\frac{1}{2}$	1	1	$\delta_{2k} < \frac{\sqrt{2}}{2}$
1	2	1	$\delta_{3k} < \sqrt{\frac{2}{3}}$
$\frac{1}{2}$	4	1	$\delta_{3k} < \sqrt{\frac{8}{9}}$
1	6	1	$\delta_{7k} < \sqrt{\frac{6}{7}}$
$\frac{1}{2}$	36	1	$\delta_{7k} < \sqrt{\frac{216}{217}}$

This theorem says that although the ℓ_q-analysis minimization is a nonconvex optimization problem with many local minimums, one still can find the global optimum under the condition (6). As pointed out by Blanchard and Thompson [2009], the higher-order RIP condition, just as (6), is easier to be satisfied by a larger subset of matrix ensemble such as Gaussian random matrices. Thus, our result is meaningful both theoretically and practically.

It is easy to verify that the right-hand side of the condition (6) is monotonically decreasing with respect to $q \in (0, 1]$ when $t \geq 1$. Therefore, in terms of the $\mathcal{A}$-RIP constant $\delta_{(t^q+1)k}$ with order more than $2k$, the condition (6) is relaxed if we use the ℓ_q-analysis minimization ($q < 1$) instead of the ℓ_1-analysis minimization. A resulted benefit is that the nonconvex ℓ_q-analysis minimization allows more sampling matrices to be used than its convex counterpart in compressed sensing. Given a ρ, a larger condition number κ will make the condition (6) more restrictive, because the value of the inequality's right-hand side becomes smaller. In other words, an analysis operator with a too large condition number could let the ℓ_q-analysis minimization fail to do recovery. This provides hints on the evaluation of the analysis operator. For example, it is reasonable to choose a tight frame as the analysis operator in some signal processing applications. When q tends to zero, the following result is straightforward.

Corollary 3 *Let $\boldsymbol{\beta} \in \mathbb{R}^d$, $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$, and $||\mathbf{D}\boldsymbol{\beta}||_0 = k$. Assume that $\delta_{2k} < \frac{\rho(1-\kappa^4)+\kappa^2\sqrt{4\rho+1}}{\rho(\kappa^2+1)^2+\kappa^2}$ with $\rho = t^{-2}/4$. Then there is some small enough $q > 0$ such that the minimizer of the ℓ_q-analysis minimization problem (4) with $\epsilon = 0$ is exactly $\boldsymbol{\beta}$.*

Remark 1 In the case $\mathbf{D} = \mathbf{I}$ and $q = 1$, the condition (6) is the same as the one of Theorem 1.1 in [Cai and Zhang, 2014] which is a sharp condition for the basis pursuit problem. Table 3.1 shows several sufficient conditions for exact recovery via the ℓ_q-analysis optimization. Compared to previous work, our results promote a significant improvement. For example, for the ℓ_1-analysis minimization, our condition $\delta_{2k} < \frac{\sqrt{2}}{2}$ is weaker than the conditions $\delta_{2k} < 0.08$ in [Candés et al., 2010], $\delta_{2k} < 0.2$ in [Liu et al., 2012], $\delta_{2k} < 0.47$ in [Lin et al.,

2013], $\delta_{2k} < 0.49$ in [Li and Lin, 2014]; and our $\delta_{7k} < \sqrt{\frac{6}{7}}$ is weaker than $\delta_{7k} < 0.6$ in [Candés et al., 2010] and [Aldroubi et al., 2012], $\delta_{7k} < 0.687$ in [Lin et al., 2013]. While for the ℓ_q-analysis minimization ($q < 1$), the D-RIP conditions $\delta_{7k} < \frac{6-3(2/3)^{2/q-2}}{6-(2/3)^{2/q-2}}$ [Aldroubi et al., 2012] and $\delta_{2k} < 0.5$ [Li and Lin, 2014] are both stronger than our condition (6). Note that above results all consider the Parseval tight frame case ($\kappa = 1$).

3.2 Stable Recovery in Noisy Case

Now we consider the case that the observation is contaminated with stochastic noise ($\epsilon \neq 0$) and the analysis vector $\mathbf{D}\boldsymbol{\beta}^*$ is approximately sparse. This is of great interest for many applications. Our goal is to provide estimate error bound between the population parameter $\boldsymbol{\beta}^*$ and the minimizer $\hat{\boldsymbol{\beta}}$ of the ℓ_q-analysis minimization (4).

Theorem 3 *Let $\boldsymbol{\beta}^* \in \mathbb{R}^d$, $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \mathbf{w}$, and $||\mathbf{w}|| \leq \epsilon$. Assume that $\mathbf{D} \in \mathbb{R}^{n\times d}$ has full column rank, and its condition number is upper bounded by $\kappa < \sqrt{\frac{2\rho+1+\sqrt{4\rho+1}}{2\rho}}$. If $\mathbf{X} \in \mathbb{R}^{m\times d}$ satisfies the $\mathcal{A}$-RIP over the set $\mathcal{A} = \{\mathbf{Dv} : ||\mathbf{v}||_0 \leq (t^q+1)k\}$ with $k, t^q k \in \mathbb{N}, t > 0, q \in (0,1]$, i.e.,*

$$\delta_{(t^q+1)k} < \frac{\rho(1-\kappa^4) + \kappa^2\sqrt{4\rho+1}}{\rho(\kappa^2+1)^2 + \kappa^2} \tag{7}$$

with $\rho = 4^{1/q-2}t^{q-2}$, then the minimizer $\hat{\boldsymbol{\beta}}$ of the ℓ_q-analysis minimization problem (4) obeys

$$||\boldsymbol{D}\hat{\boldsymbol{\beta}} - \boldsymbol{D}\boldsymbol{\beta}^*||_q^q \leq 2c_1^q k^{1-q/2}\epsilon^q + 2(2c_2^q+1)\sigma_k(\mathbf{D}\boldsymbol{\beta}^*)_q^q,$$
$$||\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*||_2 \leq \frac{2c_1}{\sigma_{min}(\boldsymbol{D})}\epsilon + \frac{2^{1/q}(2c_2+1)}{\sigma_{min}(\boldsymbol{D})}\frac{\sigma_k(\mathbf{D}\boldsymbol{\beta}^*)_q}{k^{1/q-1/2}},$$

where

$$\begin{aligned}
c_0 &= (\frac{1}{2}-\mu)^2(1+\delta_{(t^q+1)k})\kappa^2 - \frac{1}{4}(1-\delta_{(t^q+1)k}) \\
&\quad + \rho\mu^2(\kappa^2(1+\delta_{(t^q+1)k}) - (1-\delta_{(t^q+1)k})), \\
c_1 &= \frac{2\kappa(\mu-\mu^2)\sqrt{1+\delta_{(t^q+1)k}}\sigma_{max}(\mathbf{D})}{-c_0}, \\
c_2 &= \frac{2\rho\mu^2(\kappa^2(1+\delta_{(t^q+1)k}) - (1-\delta_{(t^q+1)k}))}{-c_0} \\
&\quad + \frac{\sqrt{-c_0\rho\mu^2(\kappa^2(1+\delta_{(t^q+1)k}) - (1-\delta_{(t^q+1)k}))}}{-c_0},
\end{aligned}$$

and $\mu > 0$ is a constant depending on ρ and κ.

The ℓ_2 error bound shows that the ℓ_q-analysis optimization can stably recover the approximately cosparse vector in presence of noise. Again, we can see that a too ill-conditioned analysis operator leads to bad performance. Additionally, a ℓ_q error bound of the difference between $\boldsymbol{\beta}^*$ and $\hat{\boldsymbol{\beta}}$ in the analysis domain is provided, which will be used to show the advantage of the ℓ_q-analysis minimization in the next subsection.

The linear model (1) with Gaussian noise is of particular interest in machine learning and signal processing. Lemma 1 in [Cai et al., 2009b] shows that the noise vector $\mathbf{w} \sim N(0, \sigma^2\mathbf{I})$ is upper bounded by $\sigma\sqrt{m + 2\sqrt{m \log m}}$. The following result is thus evident.

Corollary 4 *If the matrix $\mathbf{X} \in \mathbb{R}^{m\times d}$ satisfies the $\mathcal{A}$-RIP condition (7) and the noise vector $\mathbf{w} \sim N(0, \sigma^2\mathbf{I})$, then the minimizer $\hat{\boldsymbol{\beta}}$ of (4) satisfies*

$$||\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*||_2 \leq \frac{2c_1}{\sigma_{min}(\boldsymbol{D})}\sigma\sqrt{m + 2\sqrt{m \log m}} + \frac{2^{1/q}(2c_2 + 1)}{\sigma_{min}(\boldsymbol{D})}\frac{\sigma_k(\mathbf{D}\boldsymbol{\beta})_q}{k^{1/q-1/2}}$$

with probability at least $1 - \frac{1}{m}$.

3.3 Benefits of Nonconvex ℓ_q-analysis Minimization

The advantage of the nonconvex ℓ_q-analysis minimization over its convex counterpart is two-fold: the nonconvex approach can do recovery with a lower sample complexity and in a wider range of cosparsity.

The following theorem is a natural extension of Theorem 2.7 of Fourcart et al. [2010] in which $\mathbf{D} = \mathbf{I}$.

Theorem 4 *Let $m, n, k \in \mathbb{N}$ with $m, k < n$. Suppose that a matrix $\mathbf{X} \in \mathbb{R}^{m\times d}$, a linear operator $\mathbf{D} \in \mathbb{R}^{n\times d}$ and a decoder $\triangle : \mathbb{R}^m \to \mathbb{R}^d$ solving $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$ satisfy for all $\boldsymbol{\beta} \in \mathbb{R}^d$,*

$$||\mathbf{D}\boldsymbol{\beta} - \triangle(\mathbf{X}\boldsymbol{\beta})||_q^q \leq C\sigma_k(\mathbf{D}\boldsymbol{\beta})_q^q$$

with some constant $C > 0$ and some $0 < q \leq 1$. Then the minimal number of samples m obeys

$$m \geq C_1 qk \log(n/4k)$$

with $k = ||\mathbf{D}\boldsymbol{\beta}||_0$ and $C_1 = 1/(2\log(2C + 3))$.

Define the decoder $\triangle(\mathbf{X}\boldsymbol{\beta}) := \mathbf{D}\triangle_0(\mathbf{y})$ with $\triangle_0(\mathbf{y}) := \operatorname{argmin}_{\boldsymbol{\beta}, \mathbf{y}=\mathbf{X}\boldsymbol{\beta}} ||\mathbf{D}\boldsymbol{\beta}||_q^q$. Combining with the ℓ_q error bound in Theorem 3, we attain the following result.

Corollary 5 *To recover the population parameter $\boldsymbol{\beta}^*$, the minimal number of samples m for the ℓ_q-analysis minimization must obey*

$$m \geq C_2 qk \log(n/4k),$$

where $k = ||\mathbf{D}\boldsymbol{\beta}^*||_0$ and $C_2 = 1/(2\log(8c_2^q + 7))$ (c_2 is the constant in Theorem 3).

Remark 2 In our analysis of the estimate error above, we used the $\mathcal{A}$-RIP over the set $\mathcal{A} = \{\mathbf{Dv} : ||\mathbf{v}||_0 \leq (t^q + 1)k\}$, i.e., the D-RIP. As pointed out by Candés et al. [2010], random matrices with Gaussian, subgaussian, or Bernoulli entries satisfy the D-RIP with sample complexity on the order of $k \log(n/k)$. It is consistent with Corollary 5 in the case $q = 1$. However, we see that the ℓ_q-analysis minimization can have a lower sample complexity than the ℓ_1-analysis minimization. Additionally, to guarantee the uniqueness of a l-cosparse solution of ℓ_0-analysis minimization, the minimal number of samples required should satisfy the following condition:

$$m \geq 2 \cdot \max_{|\Lambda| \geq l} \dim(\mathcal{W}_\Lambda), \tag{8}$$

where $\mathcal{W}_\Lambda = \text{Null}(\mathbf{D}_\Lambda)$. Please refer to Nam et al. [2011] for more details. Therefore, the sample complexity of ℓ_q-analysis minimization is lower bounded by $2 \cdot \max_{|\Lambda| \geq l} \dim(\mathcal{W}_\Lambda)$.

The condition (6) guarantees that cosparse vectors can be exactly recovered via the ℓ_q-analysis minimization. Define S_q $(0 < q \leq 1)$ as the largest value of the sparsity $S \in \mathbb{N}$ of the analysis vector $\mathbf{D}\boldsymbol{\beta}$ such that the condition (6) holds for some $t^q \in \frac{1}{S}\mathbb{N}$. The following theorem indicates the relationship between S_q with $q < 1$ and S_1 with $q = 1$.

Theorem 5 *Suppose that there exist $S_1 \in \mathbb{N}$ and $t \in \frac{1}{S_1}\mathbb{N}$ such that*

$$\delta_{(t+1)S_1} < \frac{\rho(1-\kappa^4) + \kappa^2\sqrt{4\rho+1}}{\rho(\kappa^2+1)^2 + \kappa^2}$$

with $\rho = \frac{1}{4}t^{-1}$. Then there exist $S_q \in \mathbb{N}$ and $l^q \in \frac{1}{S_q}\mathbb{N}$ obeying

$$S_q = \left\lfloor \frac{t+1}{t^{\frac{q}{2-q}} + 1} S_1 \right\rfloor \tag{9}$$

such that $(t+1)S_1 = (l^q + 1)S_q$ and

$$\delta_{(l^q+1)S_q} < \frac{\rho(1-\kappa^4) + \kappa^2\sqrt{4\rho+1}}{\rho(\kappa^2+1)^2 + \kappa^2}$$

with $\rho = \frac{1}{4}l^{q-2}$.

It can be verified that Theorem 5 also holds for the condition (7). The equation (9) states that the ℓ_q-analysis minimization with $q < 1$ can do recovery in a wider range of cosparsity than the ℓ_1-analysis minimization. For example, if $\delta_{5S_1} < \frac{2\sqrt{5}}{5}$, then the $\ell_{\frac{2}{3}}$-analysis minimization can recover a vector $\boldsymbol{\beta}$ with $||\mathbf{D}\boldsymbol{\beta}||_0 = S_{\frac{2}{3}} = \lfloor \frac{5}{3}S_1 \rfloor$.

4 Iteratively Reweighted Method for ℓ_q-analysis Minimization

The iteratively reweighted method is a classical approach to deal with the ℓ_q norm related optimization problem; see [Gorodnitsky and Rao, 1997, Chartrand and Yin, 2008, Daubechies et al., 2010, Lu, 2014]. Inspired by them, we develop an iteratively reweighted method to solve the ℓ_q-analysis optimization. We reformulate (4) as the following unconstrained optimization problem:

$$\min_{\boldsymbol{\beta}} \Big\{ \frac{1}{2}||\mathbf{y}-\mathbf{X}\boldsymbol{\beta}||_2^2 + \lambda||\mathbf{D}\boldsymbol{\beta}||_q^q \Big\}. \tag{10}$$

It is hard to solve (10) directly due to the nonsmoothness and nonseparability of the ℓ_q norm term. We provide a way to deal with the ℓ_q norm under the variational framework.

Note that the function $||\mathbf{D}\boldsymbol{\beta}||_q^q$ is concave with respect to $|\mathbf{D}\boldsymbol{\beta}|^\alpha = (|\mathbf{D}_{1\cdot}\boldsymbol{\beta}|^\alpha, \ldots, |\mathbf{D}_{n\cdot}\boldsymbol{\beta}|^\alpha)^T$ for $\alpha \geq 1$. Thus there exists a variational upper bound of $||\mathbf{D}\boldsymbol{\beta}||_q^q$. Given a positive vector $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_n)^T$, we have the following variational formulation,

$$\begin{aligned} ||\mathbf{D}\boldsymbol{\beta}||_q^q &= \sum_{i=1}^n (|\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^\alpha)^{\frac{q}{\alpha}} \\ &= \min_{\boldsymbol{\eta}>\mathbf{0}} \Big\{ J_\alpha \triangleq \frac{q}{\alpha} \sum_{i=1}^n \Big(\eta_i |\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^\alpha + \frac{\alpha-q}{q} \frac{1}{\eta_i^{\frac{q}{\alpha-q}}} \Big) \Big\} \end{aligned}$$

for $\alpha \geq 1$ and $0 < q \leq 1$. The function J_α is jointly convex in $(\boldsymbol{\beta}, \boldsymbol{\eta})$. Its minimum is achieved at $\eta_i = 1/|\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^{\alpha-q}$, $i = 1, \ldots, n$. However, when $\boldsymbol{\beta}$ is orthogonal to some $\mathbf{D}_{i\cdot}$, the weight vector $\boldsymbol{\eta}$ may include infinite components. To avoid an infinite weight, we add a smoothing term $q/\alpha \sum_{i=1}^n \eta_i \varepsilon^\alpha$ ($\varepsilon \geq 0$) to J_α.

Using the above variational formulation, we obtain an approximation of the problem (10) as

$$\begin{aligned} \min_{\boldsymbol{\beta}} \Big\{ F(\boldsymbol{\beta}, \varepsilon) \triangleq \min_{\boldsymbol{\eta}>\mathbf{0}} & \frac{1}{2}||\mathbf{y}-\mathbf{X}\boldsymbol{\beta}||_2^2 \\ & + \frac{\lambda q}{\alpha} \sum_{i=1}^n \Big[\eta_i (|\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^\alpha + \varepsilon^\alpha) + \frac{\alpha-q}{q} \frac{1}{\eta_i^{\frac{q}{\alpha-q}}} \Big] \Big\}. \end{aligned} \tag{11}$$

We then develop an alternating minimization algorithm, which consists of three steps. The first step calculates $\boldsymbol{\eta}$ with $\boldsymbol{\beta}$ fixed via

$$\boldsymbol{\eta}^{(k)} = \operatorname*{argmin}_{\boldsymbol{\eta}>\mathbf{0}} \Big\{ \sum_{i=1}^n \Big[\eta_i (|\mathbf{D}_{i\cdot}\boldsymbol{\beta}^{(k-1)}|^\alpha + \varepsilon^\alpha) + \frac{\alpha-q}{q} \frac{1}{\eta_i^{\frac{q}{\alpha-q}}} \Big] \Big\},$$

which has a closed form solution. The second step calculates $\boldsymbol{\beta}$ with $\boldsymbol{\eta}$ fixed via

$$\boldsymbol{\beta}^{(k)} = \operatorname*{argmin}_{\boldsymbol{\beta}\in\mathbb{R}^d} \Big\{ \frac{1}{2}||\mathbf{y}-\mathbf{X}\boldsymbol{\beta}||_2^2 + \frac{\lambda q}{\alpha} \sum_{i=1}^n \eta_i^{(k)} |\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^\alpha \Big\},$$

which is a weighted ℓ_α-minimization problem. Particularly, the case $\alpha = 2$ corresponds to a least squares problem which can be solved efficiently. The third step updates the smoothing parameter ε according to the following rule [3]

$$\varepsilon^{(k)} = \min\{\varepsilon^{(k-1)}, \rho \cdot r(\mathbf{D}\boldsymbol{\beta}^{(k)})_l\} \quad \text{with } \rho \in (0,1),$$

where $r(\mathbf{D}\boldsymbol{\beta})_l$ is the l-th smallest element of the set $\{|\mathbf{D}_{j\cdot}\boldsymbol{\beta}| : j = 1, \ldots, n\}$. $\boldsymbol{\beta}$ is a l-cosparse vector if and only if $r(\mathbf{D}\boldsymbol{\beta})_l = 0$. The algorithm stops when $\varepsilon = 0$.

Algorithm 1 The CoIRLq Algorithm

Input: l, $\mathbf{X}$, $\mathbf{y}$, $\mathbf{D} = [\mathbf{D}_{1\cdot}; \ldots; \mathbf{D}_{n\cdot}]$.
Init: Choose $\boldsymbol{\beta}^{(0)}$ such that $\mathbf{X}\boldsymbol{\beta}^{(0)} = \mathbf{y}$ and $\varepsilon^{(0)} = 1$.
while $\|\boldsymbol{\beta}^{(k+1)} - \boldsymbol{\beta}^{(k)}\|_\infty > \tau$ or $\varepsilon^{(k)} \neq 0$ **do**
Update

$$\eta_i^{(k)} = \left(|\mathbf{D}_{i\cdot}\boldsymbol{\beta}^{(k-1)}|^\alpha + (\varepsilon^{(k-1)})^\alpha\right)^{\frac{q}{\alpha}-1}, i = 1, \ldots, n.$$

Update

$$\boldsymbol{\beta}^{(k)} = \underset{\boldsymbol{\beta}\in\mathbb{R}^d}{\operatorname{argmin}} \Big\{\frac{1}{2}||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}||_2^2 + \frac{\lambda q}{\alpha}\sum_{i=1}^{n} \eta_i^{(k)}|\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^\alpha\Big\}.$$

Update

$$\varepsilon^{(k)} = \min\{\varepsilon^{(k-1)}, \rho \cdot r(\mathbf{D}\boldsymbol{\beta}^{(k)})_l\} \text{ with } \rho \in (0,1).$$

end while
Output: $\boldsymbol{\beta}$

4.1 Convergence Analysis

Our analysis is based on the optimization problem (11) with the objective function $F(\boldsymbol{\beta}, \varepsilon)$. Noting that $\boldsymbol{\eta}^{(k+1)}$ is a function of $\boldsymbol{\beta}^{(k)}$ and $\varepsilon^{(k)}$, we define the following objective function

$$Q(\boldsymbol{\beta}, \varepsilon|\boldsymbol{\beta}^{(k)}, \varepsilon^{(k)}) \triangleq \frac{1}{2}||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}||_2^2 + \frac{\lambda q}{\alpha}\sum_{i=1}^{n}\Big[\eta_i^{(k+1)}(|\mathbf{D}_{i\cdot}\boldsymbol{\beta}|^\alpha + \varepsilon^\alpha) + \frac{\alpha - q}{q}\frac{1}{\eta_i^{(k+1)\frac{q}{\alpha-q}}}\Big].$$

Lemma 1 *Assume that the analysis operator* $\mathbf{D}$ *has full column rank. Let* $\{(\boldsymbol{\beta}^{(k)}, \varepsilon^{(k)}) : k = 1, 2, \ldots\}$ *be a sequence generated by the CoIRLq algorithm. Then,*

$$||\boldsymbol{\beta}^{(k)}||_2 \leq \sigma_{min}^{-1}(\mathbf{D})(F(\boldsymbol{\beta}^{(0)}, \varepsilon^{(0)})/\lambda)^{1/q}$$

[3] Various strategies can be applied to update ε. For example, we can keep ε as a small fixed value. It is preferred to choose a sequence of $\{\varepsilon^{(k)}\}$ tending to zero [Daubechies et al., 2010].

and

$$F(\boldsymbol{\beta}^{(k+1)}, \varepsilon^{(k+1)}) \leq F(\boldsymbol{\beta}^{(k)}, \varepsilon^{(k)}),$$

with equality holding if and only if $\boldsymbol{\beta}^{(k+1)} = \boldsymbol{\beta}^{(k)}$ *and* $\varepsilon^{(k+1)} = \varepsilon^{(k)}$.

The boundedness of $||\boldsymbol{\beta}^{(k)}||_2$ implies that the sequence $\{\boldsymbol{\beta}^{(k)}\}$ converges to some accumulation point. We can immediately derive the convergence property of the CoIRLq algorithm from Zangwill's *global convergence theorem* or the literature [Sriperumbudur and Lanckriet, 2009]. Here we omit the detail. Finally, it is easy to verify that when $\varepsilon^* = 0$, $\boldsymbol{\beta}^*$ is a stationary point of (10).

4.2 Recovery Guarantee Analysis

To uniquely recover the true parameter, the linear operator $\mathbf{X} : \mathcal{A} \to \mathbb{R}^m$ must be a one-to-one map. Define a set $\bar{\mathcal{A}} = \{\boldsymbol{\beta} = \boldsymbol{\beta}_1 + \boldsymbol{\beta}_2 : \boldsymbol{\beta}_1, \boldsymbol{\beta}_2 \in \mathcal{A}\}$. Blumensath and Davies [2008] showed that a necessary condition for the existence of a one-to-one map requires that $\delta_{\bar{\mathcal{A}}} < 1$ ($\delta_{\mathcal{A}} \leq \delta_{\bar{\mathcal{A}}}$). For any two l-cosparse vectors $\boldsymbol{\beta}_1, \boldsymbol{\beta}_2 \in \mathcal{A} = \{\boldsymbol{\beta} : \mathbf{D}_\Lambda \boldsymbol{\beta} = \mathbf{0}, |\Lambda| \geq l\}$, denote $T_1 = supp(\mathbf{D}\boldsymbol{\beta}_1)$, $T_2 = supp(\mathbf{D}\boldsymbol{\beta}_2)$, $\Lambda_1 = cosupp(\mathbf{D}\boldsymbol{\beta}_1)$ and $\Lambda_2 = cosupp(\mathbf{D}\boldsymbol{\beta}_2)$. Since $supp(\mathbf{D}(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_2)) \subseteq T_1 \cup T_2$, we have $cosupp(\mathbf{D}(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_2)) \supseteq (T_1 \cup T_2)^c = T_1^c \cap T_2^c = \Lambda_1 \cap \Lambda_2$. Moreover, we also have $|\Lambda_1 \cap \Lambda_2| = n - |T_1 \cup T_2| \geq n - (n-l) - (n-l) = 2l - n$. Thus it requires that the linear operator $\mathbf{X}$ satisfies the $\mathcal{A}$-RIP with $\delta_{2l-n} < 1$ to uniquely recover any l-cosparse vector from the set $\mathcal{A} = \{\boldsymbol{\beta} : \mathbf{D}_\Lambda \boldsymbol{\beta} = 0, |\Lambda| \geq l\}$. Otherwise, there would exist two l-cosparse vectors $\boldsymbol{\beta}_1 \neq \boldsymbol{\beta}_2$ such that $\mathbf{X}(\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2) = 0$. Giryes et al. [2013] showed that there exists a random matrix $\mathbf{X}$ satisfying such a requirement with high probability.

Theorem 6 *Let* $\boldsymbol{\beta}^* \in \mathbb{R}^d$ *be a* l*-cosparse vector, and* $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \mathbf{w}$ *with* $||\mathbf{w}||_2 \leq \epsilon$. *Assume that* $\mathbf{X}$ *satisfies the* $\mathcal{A}$*-RIP over the set* $\mathcal{A} = \{\boldsymbol{\beta} : \mathbf{D}_\Lambda \boldsymbol{\beta} = \mathbf{0}, |\Lambda| \geq l\}$ *of order* $2l - n$ *with* $\delta_{2l-n} < 1$. *Then the solution* $\hat{\boldsymbol{\beta}}$ *obtained by the CoIRLq algorithm obeys*

$$||\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*||_2 \leq C_1 \sqrt{\lambda} + C_2 \epsilon,$$

where C_1 *and* C_2 *are constants depending on* δ_{2l-n}.

We can see that the CoIRLq algorithm can recover an approximate solution away from the true parameter vector by a factor of $\sqrt{\lambda}$ in the noiseless case.

5 Numerical Analysis

In this section we conduct numerical analysis of the ℓ_q-analysis minimization method on both simulated data and real data, and compare the performance of the case $q < 1$ and the case $q = 1$. We set $\alpha = 2$ in the CoIRLq algorithm.

5.1 Cosparse Vector Recovery

We generate the simulated datasets according to

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{w},$$

where $\mathbf{w} \sim N(\mathbf{0}, \sigma\boldsymbol{I})$. The sampling matrix $\mathbf{X}$ is drawn independently from the normal distribution with normalized columns. The analysis operator $\mathbf{D}$ is constructed such that $\mathbf{D}^T$is a random tight frame. To generate a l-cosparse vector $\boldsymbol{\beta}$, we first choose l rows randomly from $\mathbf{D}$ and form $\mathbf{D}_\Lambda$.Then we generate a vector which lies in the null space of $\mathbf{D}_\Lambda$. The recovery is deemed to be successful if the recovery relative error $||\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*||_2/||\boldsymbol{\beta}^*||_2 \leq 1e-4$.

In the first experiment, we test the vector recovery capability of the CoIRLq method with $q = 0.7$. We set $m = 80, n = 144, d = 120, l = 99$, and $\sigma = 0$. Figure 1 illustrates that the CoIRLq method recovers the original vector perfectly.

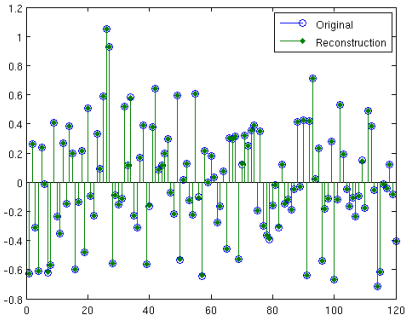
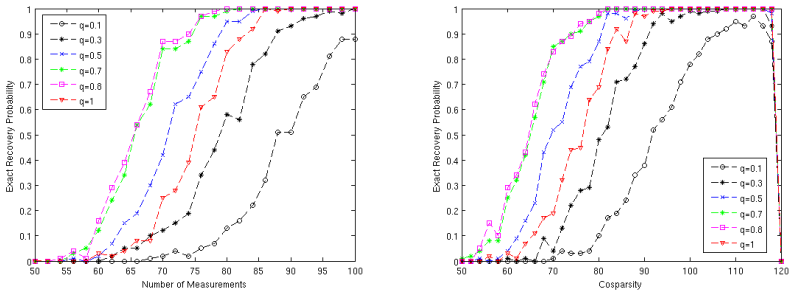


Figure 1: Cosparse vector recovery.

In the second experiment, we test the CoIRLq method on a range of sample size and cosparsity with different q. Although the optimal tuning parameter λ depends on q, a small enough λ is able to ensure that $\mathbf{y}$ approximately equals to $\mathbf{X}\boldsymbol{\beta}$ in the noiseless case. Thus, we set $\lambda = 1e-4$ for all q and $\sigma = 0$. Figure 2 reports the result with 100 repetitions on every dataset. We can see that the CoIRLq method with $q = 0.5, 0.7, 0.8$ can achieve exact recovery in a wider range of cosparsity and with fewer samples than with $q = 1$. In addition, it should be noted that small $q = 0.1$ or $q = 0.3$ do not perform better than relatively large $q = 0.7, 0.8$, because a too small q leads to a hard-solving problem. Note that there is a drop of recovery probability where the cosparsity $l = 118$ [4]. This is because it is hard to algorithmically recover a vector residing in a subspace with a small dimension; please also refer to [Nam et al., 2011].

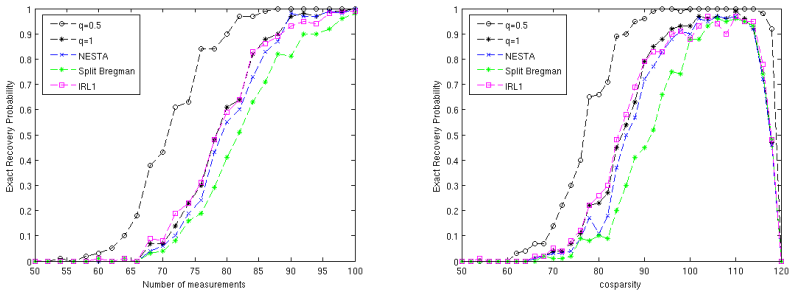
In the third experiment, we compare the CoIRLq method with three state-of-the-art methods for the ℓ_1-analysis minimization problem including NESTA

[4]When $l = 120$, a zero vector is generated by our codes. So the recovery probability in cosparsity $l = 120$ is zero.



$n = 144,\ d = 120,\ l = 99$ $m = 90,\ n = 144,\ d = 120$

Figure 2: Exact recovery probability of the CoIRLq method.



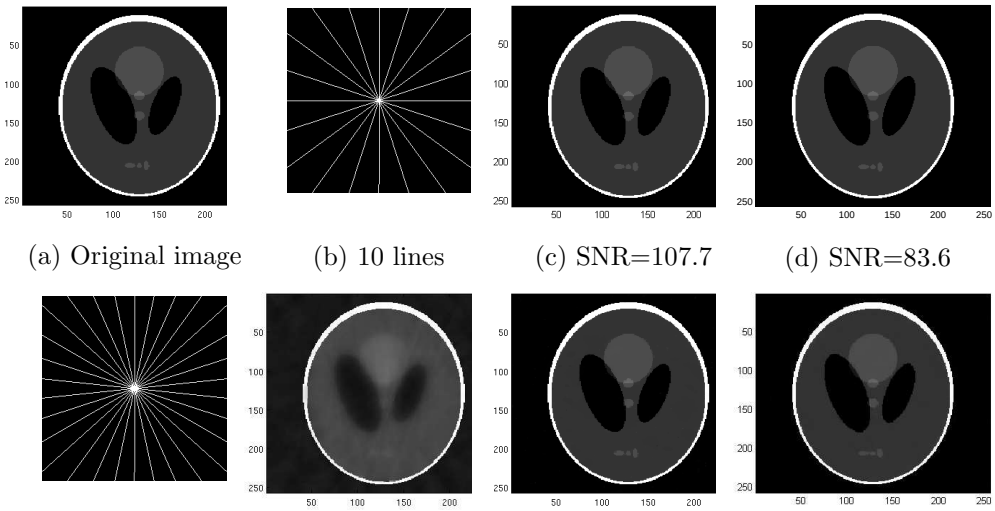
$n = 144,\ d = 120,\ l = 99$ $m = 90,\ n = 144,\ d = 120$

Figure 3: Recovery probability of the CoIRLq, NESTA, IRL1 and split Bregman methods.

(http://statweb.stanford.edu/∼candes/nesta/), split Bregman method, and iteratively reweighted ℓ_1 (IRL1) method. Set the noise level $\sigma = 0.01$. The parameter λ is tuned via the grid search method. We run these methods in a range of sample size and cosparsity. Figure 3 reports the result with 100 repetitions on every dataset. We can see that the nonconvex ℓ_q-analysis minimization with $q < 1$ is more capable of achieving exact recovery against noise than the convex ℓ_1-analysis minimization. Moreover, the nonconvex approach can obtain exact recovery with fewer samples or in a wider range of cosparsity than the convex counterpart. Moreover, we found that the CoIRLq algorithm in the case $q < 1$ often needs less iterations than in the case $q = 1$.

5.2 Image Restoration Experiment

In this section we demonstrate the effectiveness of the ℓ_q-analysis minimization on the Shepp Logan phantom reconstruction problem. In computed tomography, an image can not be observed directly. Instead, we can only obtain its 2D Fourier transform coefficients along a few radial lines due to certain limitations. This sampling process can be modeled as a measurement matrix $\mathbf{X}$. The goal



(a) Original image (b) 10 lines (c) SNR=107.7 (d) SNR=83.6

(e) 15 lines (f) SNR=30.5 (g) SNR=45.7 (h) SNR=43

Figure 4: (a) Original Shepp Logan Phantom image; (b) Sampling locations along 10 radial lines; (c) Exact reconstruction via CoIRLq (q=0.7) with 10 lines without noise; (d) Exact reconstruction via CoIRLq (q=1) with 12 lines without noise; (e) Sampling locations along 15 radial lines; (f) Reconstruction via GAP with 15 lines and noise level $\sigma = 0.01$; (g) Reconstruction via CoIRLq (q=0.7) with 15 lines and noise level $\sigma = 0.01$; (h) Reconstruction via CoIRLq (q=1) with 15 lines and noise level $\sigma = 0.01$.

is to reconstruct the image from the observation.

The experimental program is set as follows. The image dimension is of 256×256, namely $d = 65536$. The measurement matrix $\mathbf{X}$ is a two dimensional Fourier transform which measures the image's Fourier transform along a few radial lines. The analysis operator is a finite difference operator $\mathbf{D}_{\text{2D-DIF}}$ whose size is roughly twice the image size, namely $n = 130560$. Since the number of nonzero analysis coefficients is $n-l = 2546$, the cosparsity used is $l = n-2546 = 128014$. The number of measurements depends on the number of radial lines used. To show the reconstruction capability of the CoIRLq method, we conduct the following experiments (the parameter λ is tuned via grid search). First, we compare our method with the greedy analysis pursuit (GAP http://www.small-project.eu/software-data) method for the ℓ_0-analysis minimization.

Figures 4-(f), (g) and (h) show that our method performs better than the GAP method in the noisy case. We can see that the CoIRLq method with $q < 1$ is more robust to noise than the case with $q = 1$. Second, we take an experiment using 10 radial lines without noise. The corresponding number of measurements is $m = 2282$, which is approximately 3.48% of the image size. Figure 4-(c) demonstrates that the CoIRLq ($q = 0.7$) method with 10 lines obtains perfect

reconstruction. Figure 4-(d) shows that the CoIRLq ($q = 1$) method with 12 lines attains perfect reconstruction. However, the GAP method needs at least 12 radial lines to achieve exact recovery; see [Nam et al., 2011].

6 Conclusion

In this paper we have conducted the theoretical analysis and developed the computational method, for the ℓ_q-analysis minimization problem. Theoretically, we have established weaker conditions for exact recover in noiseless case and a tighter non-asymptotic upper bound of estimate error in noisy case. In particular, we have presented a necessary and sufficient condition guaranteeing exact recovery. Additionally, we have shown that the nonconvex ℓ_q-analysis optimization can do recovery with a lower sample complexity and in a wider range of cosparsity. Computationally, we have devised an iteratively reweighted method to solve the ℓ_q-analysis optimization problem. Empirical results have illustrated that our iteratively reweighted method outperforms the state-of-the-art methods.

References

A. Aldroubi, X. Chen, and A. Powell. Perturbations of measurement matrices and dictionaries in compressed sensing. *Applied and Computational Harmonic Analysis*, 33(2):282–291, 2012.

S. Becker, J. Bobin, and E. J. Candés. Nesta: A fast and accurate first-order method for sparse recovery. *SIAM Journal on Imaging Sciences*, 4(1):1–39, 2011.

J. D. Blanchard and A. Thompson. On support sizes of restricted isometry constants. *Applied and Computational Harmonic Analysis*, 29(3):382–390, 2009.

T. Blumensath and M. E. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Transactions on Information Theory*, 55(4):1872–1882, 2008.

J. F. Cai, S. Osher, and Z. W. Shen. Split bregman methods and frame based image restoration. *Multiscale Modeling and Simulation*, 8(2):337–369, 2009a.

T. T. Cai and A. Zhang. Sparse representation of a polytope and recovery of sparse signals and low-rank matrices. *IEEE Transactions on Information Theory*, 60(1):122–132, 2014.

T. T. Cai, G. Xu, and J. Zhang. On recovery of sparse signal via ℓ_1 minimization. *IEEE Transactions on Information Theory*, 55(7):3388–3397, 2009b.

E. J. Candés and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51:4203–4215, 2004.

E. J. Candés, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted l1 minimization. *J. Fourier Anal. Appl.*, 14:877–905, 2007.

E. J. Candés, Y. C. Eldar, D. Needle, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2010.

A. Chambolle and P. Lions. Image recovery via total variation minimization and related problems. *Numer. Math.*, 76(2):167–188, 1997.

R. Chartrand and W. Yin. Iteratively reweighted algorithms for compressive sensing. In *33rd International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2008.

S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20:33–61, 1998.

A. Cohen, W. Dahmen, and R. Devore. Compressed sensing and best k-term approximation. *J. Amer. Math. Soc*, pages 211–231, 2009.

I. Daubechies, R. Devore, M. Fornasier, and C. S. Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 2010.

J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of American Statistical Association*, pages 947–968, 2001.

S. Fourcart, A. Pajor, H. Rauhut, and T. Ullrich. The gelfand widths of ℓ_p-balls for $0 < p \leq 1$. *Journal of Complexity*, 26(6):629–640, 2010.

R. Giryes, S. Nam, M. Elad, R. Gribonval, and M. E. Davies. Greedy-like algorithms for the cosparse analysis model. In *Special Issum in LAA on Sparse Approximation Solutions of Linear Systems*. 2013.

I. F. Gorodnitsky and B. D. Rao. Sparse signal reconstruction from limited data using focuss: a reweighted minimum norm algorithm. *IEEE Transactions on Signal Processing*, 45(3):600–616, 1997.

S. J. Kim, K. Koh, S. Boyd, and D. Gorinevsky. ℓ_1 trend filtering. *SIAM Review*, 51(2):339–360, 2009.

S. Li and J. Lin. Compressed sensing with coherent tight frames via ℓ_q-minimization for $0 < q \leq 1$. *Inverse Problems and Imaging*, 2014.

J. Lin, S. Li, and Y. Shen. New bounds for restricted isometry constants with coherent tight frames. *IEEE Transactions on Information Theory*, 61(3): 611–621, 2013.

J. Liu, L. Yuan, and J. P. Ye. Guaranteed sparse recovery under linear transformation. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

Y. Liu, T. Mi, and S. Li. Compressed sensing with general frames via optimal-dual-based ℓ_1-analysis. *IEEE Transactions on Information Theory*, 58(7): 4201–4214, 2012.

Z. Lu. Iterative reweighted minimization methods for ℓ_p regularized unconstrained nonlinear programming. *Mathematical Programming*, 147(1-22):277–307, 2014.

M. Lysaker, A. Lundervold, and X. C. Tai. Noise removal using fourth-order partial differential equation with apllication to medical magnetic resonance images in space and time. *IEEE Transactions on Image Processing*, 41(12): 3397–3415, 2003.

S. Nam, M. E. Davies, M. Elad, and R. Gribonval. The cosparse analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2011.

D. Needell and R. Ward. Stable image reconstruction using total variation minimization. *SIAM Journal on Imaging Sciences*, 6(2):1035–1058, 2013.

L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

J. Sharpnack, A. Rinaldo, and A. Singh. Sparsistency of the edgelasso over graphs. In *AISTAT*, 2012.

B. K. Sriperumbudur and G. R. G. Lanckriet. On the convergence of the concave-convex procedure. In *Advances in Neural Information Processing Systems*, 2009.

R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc. B*, 58(1):267–288, 1996.

R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via fused lasso. *J. Royal. Statist. Soc. B*, 67(Part 1):91–108, 2005.

R. J. Tibshirani and J. Taylor. The solution path of the generalized lasso. *The Annals of Statistics*, 39(3):1335–1371, 2011.

S. Vaiter, G. Payre, C. Dossal, and J. Fadili. Robust sparse analysis regularization. *IEEE Transactions on Information Theory*, 39(3):1335–1371, 2013.

V. Viallon, S. Lambert-Lacroix, H. Hoefling, and F. Picard. On the robustness of the generalized fused lasso to prior specifications. *Statistics and Computing*, pages 1–17, 2014.